

Architectural Constructs of AMPEX DST

Clay Johnson

Ampex Systems Corporation
2345 Vantage Drive
Colorado Springs, CO 80919
Phone (719)590-9758
Fax (719)590-7526
Clay_Johnson@ampex.com

Abstract

The DST™ 800¹ automated library is a high performance, automated tape storage system, developed by AMPEX, providing mass storage to host systems.

Physical Volume Manager (PVM) is a volume server which supports either a DST 800, DST 600 stand alone tape drive or combination of DST 800 and DST 600 subsystems. The objective of the Physical Volume Manager is to provide the foundation support to allow automated and operator assisted access to the DST cartridges with continuous operation. A second objective is to create a data base about the media, its location and its usage so that the quality and utilization of the media on which specific data is recorded and the performance of the storage system may be managed.

The DST Tape Drive architecture and media provides several unique functions that enhance the ability to achieve high media space utilization and fast access. Access times are enhanced through the implementation of multiple areas (called System Zones) on the media where the media may be unloaded. This reduces positioning time in loading and unloading the cartridge. Access times are also reduced through high speed positioning in excess of 800 megabytes per second.

A DST cartridge can be partitioned into fixed size units which can be reclaimed for rewriting without invalidating other recorded data on the tape cartridge. Most tape management systems achieve space reclamation by deleting an entire tape volume, then allowing users to request a "scratch tape" or "non-specific" volume when they wish to record data to tape. Physical cartridge sizes of 25, 75, or 165 gigabytes will make this existing process inefficient or unusable. The DST cartridge partitioning capability provides an efficient mechanism for addressing the tape space utilization problem.

Overview

This paper will provide an architectural overview of the DD-2 (Ampex DST™) magnetic tape storage subsystem. The areas discussed within this paper will cover the functionality of the tape formatting and tape drive design, automated library, the device specific software, and physical volume management software. The DST library subsystem will then be discussed in the context of the National Storage Laboratory environment.

The design of DD-2 transport is a novel approach in helical recording. The unit utilizes proven transport technology developed for the video industry as it's core technology. The remainder of the system is designed specifically for the data processing industry.

¹DST is a Trademark of Ampex Systems Corporation

Ampex DD-2 Core Technology

The Ampex DD-2 design meets the stringent requirements for long media life in its approach to tape handling. The design of the Ampex transport was concurrent with the SMPTE/EBU standards development for D2 Composite Video. This allowed the adoption of several cartridge features to match specific transport requirements. One example is the size and shape of the cavity for the threading apparatus, which accommodates the large direct-drive capstan hub, in addition to four large diameter threading posts.

The unit uses this large direct-coupled capstan hub similar to high performance reel-to-reel tape drives instead of the usual pinch-roller design found in most helical transports. The advantages include fast accelerations and direction reversal without tape damage, plus elimination of the scuffing and stretching problems of pinch roller systems. Since a direct drive capstan must couple to the backside of the tape, it must be introduced inside the loop extracted from the cartridge. In order to avoid a "pop up" or moving capstan and the problems of precise registration, the capstan was placed under the cartridge elevator, so that it is introduced into the threading cavity as the cartridge is lowered onto the turntables.

In order to prevent tension buildup and potential tape damage, none of the tape guides within the transport are conventional fixed posts. Air film lubricated guides are used throughout with one exception which is a precision rotating guide that is in contact with the backside of the tape.

All motors are equipped with tachometers to provide speed, direction, or position information to the servo system, including the gear motors which power the cartridge elevator and the threading apparatus. End position sensors are used only at beginning-of-tape and end-of-tape while formatting the tape. In other situations, the servo learns the limit positions of the mechanisms and subsequently applies acceleration profiles to drive them rapidly and without crash stops. This approach also permits the machine to recover from an interruption during any phase of operation without damage to the machine or the tape.

The tape transport also features a functional intermediate tape path that allows high speed searches and reading or writing of the longitudinal tracks without the tape being in contact with helical scan drum. If tape is already threaded, high speed operations are also performed without having to unthread tape from the scanner. Thread/unthread operations are not performed over user data except when the user write performance over-rides error rate performance requirements.

The DD-2 Tape Drive architecture and media provides several unique functions that enhance the ability to achieve high media space utilization and fast access. These core technology designs allow a set of advantages unique to DD-2 tape transports:

Multiple Unload Positions

Access times are improved, as compared to traditional 3480 type cartridges, through the implementation of multiple areas (called System Zones) on the media where the media may be loaded and unloaded. Full rewind is therefore unnecessary. This reduces positioning time in loading and unloading the cartridge, eliminates mechanical actions of threading and unthreading/loading and unloading over recorded data and eliminates the wear that is always inherent in any design that requires a return to beginning of tape.

Head Life

Using DST media, helical heads are warranted for 500 hours, however, experience with helical head contact time can exceed 2000 hours. Because of the system zones and the ability to move between system zones without tape loaded to the helical scanner drum, the actual head life with tape mounted on the drive may be considerable longer.

Head Replacement

DD-2 head replacement in the field requires about 0.5 hours for fault isolation, replacement, alignment checks and repair-verification of a single head. The time required

to do the same actions on the complete head set is one hour. Some competitive helical scan implementations do not provide for head replacement in the field.

Safe Time on Stopped Tape

Whenever the flow of data to or from the tape drive is interrupted, after a pre-determined period of time, the media is moved to a system zone and unloaded from the helical drum. When data is being written, this should be a rare occurrence because of the minimum 64 (maximum 128) million bytes of buffering per drive. When in retrieval mode, returning to a system zone whenever the access queue is zero should be standard practice. In this way, if the drive is needed for a different cartridge, it is available sooner and if another access is directed at the same cartridge, the average access time is not adversely impacted by positioning the tape to the nearest system zone. Half the time it will be closer to the next access and half the time it will be farther away. With this type of drive management, the cartridge may remain mounted indefinitely without exposure to the tape or head wear.

The SCSI-2 implementation offers an interface programmable option, providing the ability to override the use of system zone for thread/unthread operations. This feature is made available for use in selected sequential access applications where optimum time-line performance is more important than achieving the best overall media life or system error rate performance.

Media Usage Life

One of the major applications for DD-2 technology is its use as a storage level in a hierarchy of storage devices, a Disk/Tape hierarchy for example. As such, the number of tape load and unload cycles, thread/unthread cycles and searches may be significant. The expected usage capabilities for Ampex DD-2 media should exceed 50,000 load/unload cycles. An even larger number of tape thread/unthread cycles spread across system zones (assuming at least three system zones on a cartridge) can be expected, and up to 5,000 shuttle forward and rewind cycles. The number of end-to-end reads using incremental motion (less than 15 MB/sec) should exceed 2,000 while the number of reads of 1 Giga Byte files using incremental motion should exceed 5,000.

An operating environment of $20 \pm 2^\circ \text{C}$ with relative humidity of $50 \pm 2\%$ will provide best overall results.

Environmental Archival Stability

Assuming cartridges are always stored within the operating environment recommended above, $20 \pm 2^\circ \text{C}$ with relative humidity of $50 \pm 2\%$ non-condensing, computer room storage of over 10 years is expected. For even longer archival stability, an environment maintained at 10°C or lower and a relative humidity of 40 % non-condensing or lower should result in archival stability exceeding 15 years.

High Speed Search

DD-2 data formats include full function use of the longitudinal tracks that can be read in either the forward or reverse direction. One of these tracks contains the geometric address of each physical block of data. This track can be searched at speeds of greater than 300 inches per second, equivalent to searching user data at more than 800 megabytes per second. Another longitudinal track is automatically recorded as user data is written to tape and provides the user with the ability to address either data block or byte offset within a user file. No user action is required to cause these tracks to be written and they provide high speed search to any point in the recorded data, not just points explicitly recorded at the time of data creation.

Data Processing Design

The DD-2 design is the only high performance cartridge based helical data storage peripheral based upon modern digital platforms (D1, D2, D3) that is commercially available today. As such, DD-2 based products will set new performance and functionality benchmarks for the data storage industry.

A DD-2 cartridge can be partitioned into fixed size units which can be reclaimed for rewriting without invalidating other recorded data on the tape cartridge. Most tape management systems achieve space reclamation by deleting an entire tape volume, then allowing users to request a "scratch tape" or "non-specific" volume when they wish to record data to tape. Physical cartridge sizes common in helical recording devices will make this existing process inefficient or unusable. The DD-2 cartridge partitioning capability provides an efficient mechanism for addressing the tape space utilization problem.

DD-2 formatting provides for three levels of Reed-Solomon error correction. In addition, data is shuffled across the 32 tracks that make up a physical block, and interleaved within the physical track so that each byte of a block has maximum separation from every other byte that make up an error correction code word. Data is then recorded using an Ampex patented process called Miller Squared. This process is self clocking, DC free rate 1/2 coding process that approaches 100% probability of flagging a burst error. This has the effect of doubling the efficiency of a Reed-Solomon code by knowing where the power of the code should be applied.

The following table summarizes the error management system.

FORMAT ITEM	FORMAT DESCRIPTION
BYTES PER TRACK	48,972
USER BYTES PER TRACK	37,495
C1 DIMENSIONS	RS (228,220,8) T=4
C2 DIMENSIONS	RS (106,96,10) T=5
C3 DIMENSIONS	RS (96,86,10) T=5
CHANNEL CODE	MILLER-SQUARED (RATE 1/2)
C1-C2 PRODUCT CODE ARRAY	IN-TRACK BLOCK INTERLEAVE WITH DIMENSIONS 456 X 106 (TWO C1 WORDS BY ONE C2 WORD)
C3 CODE CROSS-TRACK INTERLEAVE DESCRIPTION	C3 CODE WORDS INTERLEAVED ACROSS A 32-TRACK PHYSICAL BLOCK
OUTER CRC ERROR DETECTION OF C1-C2-C3 FAILURE	FOUR 64 PARITY BIT CRC CODE WORDS INTERLACED OVER 32 TRACKS WHICH PROVIDE UNDETECTED ERROR PROBABILITY OF 10^{-20}
WRITE RETRY	YES
CODING OVERHEAD	28%

The complete redesign of the D-2 video recorder into the second generation data recording technology known as DD-2 provides the following characteristics:

Error Rate

Using DST cartridges, the number of bytes read per permanent read error should exceed 10^{14} , (achieved when factoring in the effect of the interleave, write retry, write bias, and the use of system zones for all thread/unthread and cartridge load/unload operations). C3 error correction is disabled during read back check when writing in order to bias the write process. Any time C2 is unable to correct the error of any one byte, a re-try is invoked.

Error Correction Code

Reed-Solomon C1 is (228,220,8), C2 is (106,96,10), C3 is (96,86,10). All three codes are always applied, therefore, data rate, capacity, search speeds, and maximum data reliability are all achieved concurrently.

Technology

Using the core D-2 technology, a new second generation product called Data D-2 (DD-2), was created.

Data Rate

Interface rate of 20 MB/sec, 15 MB/sec device sustained, even with full error protection applied.

Head Azimuth

+14.97 degrees, -15.03 degrees for adjacent tracks provides sufficient suppression of head cross talk without undue loss of signal strength.

Data Modulation

Ampex Patented Miller Squared Code, (Self clocking, DC Free, Rate 1/2 Code).

Tape Tension

An appropriate tape tension (3+ oz.) is maintained by the tension arm and the servo system to insure stager free wraps and precise tracking registration.

Physical Format

DD-2 uses a 19 mm wide tape media. The track angle is 6.13 degrees yielding a track length of 150.8 mm. Head rotation speed is 100 RPS, four read and four write heads are mounted on the scanner, 37,495 user data bytes are recorded on each track yielding a sustained data rate of 14.998 MB/sec. Track pitch is 0.0395 mm which converts to about 69 six inch tracks per linear inch. This dense track format and rotation speed is what achieves the significant volumetric efficiency of the DD-2 design and it's superior data rate.

Data Capacity

Data capacity of the DD-2 media are 25, 75 and 165 Giga Bytes for the Small, Medium and Large size cartridges.

Data Compression

Current implementation does not provide data compression. Current system implementations tend to provide data compaction and data compression at the source of the data to gain the benefits of improved capacity and data rate throughout the system. The DD-2 format standards proposal does include provision for data compaction and compression, should it become a requirement.

Drive Configuration

DD-2 is configured with a one-to-one, transport to control unit relationship. Helical tape devices with very high data rates are streaming devices which achieve best results with large data transfers. DD-2 control units have at least 64 Mega Bytes of high speed memory buffering for each tape drive, ensuring high probability of efficient usage of the device and the channels to which they are attached.

Interface

Current interface capabilities include IPI-3 and SCSI-II available from Ampex, HIPPI (available from Maximum Strategy), ESCON, FDDI and others via an IBM RS/6000 attachment and others such as FCS and ATM fabrics as they become available.

DST 800™ ACL

DD-2 can be acquired with a varied array of automation devices. Ampex provides a 6.4 Terabyte Library (DST 800) and two other offerings are provided by other vendors, the first configurable from 3.5 to 25 Terabytes and the second configurable from 10 to 10,000 Terabytes. The DST 800 is a high performance automated cartridge library storage system utilizing an Ampex designed cartridge handling system in conjunction with a special design of Ampex's 19 mm DST helical scan tape drives. The DST 800 is designed to provide fast access to 6.4 Terabytes of on-line data within a footprint of 21 square feet.

The DST 800 utilizes the 25 gigabyte cartridge which has been specifically designed for automation. The cartridge is equipped with an indentation the cartridge picker can access to withdraw the medium from the storage bin. This eliminates the need for the complexities of a

gripping mechanism and all the moving parts associated with gripping the cartridge. This design facilitates the high performance of the operation. The cartridge actually reaches speeds of 50 inches per second while being withdrawn from its storage bin. The cartridge is drawn into a holder that is then accelerated to speeds of 90 inches per second moving between storage bins and the drive.

The DST 800 comes in two versions, Version 1 with physical volume manager (Figure 1) or Version 2 without physical volume manager (Figure 2).

Version 1:

The interface to the ACL is through an Application Programming Interface (API) Figure 1. This interface supports a set of APIs in several categories including administrative, error handling, general information and operation, media management, mount and dismount, resource reservation and operator interface. Over this interface, a tape is identified by a cartridge ID or a partition ID. The cartridge ID is the ID for the cartridge. The partition ID is either the entire cartridge or the ID of a partition on the cartridge depending the cartridge format. A cartridge can be formatted either as a one partition cartridge or with multiple partitions on the same cartridge. Each partition acting as a logically contiguous independent storage space.

The DST 800 system delivered and installed at the National Storage Laboratory is as described in version 1 above. The File Management Software is the NSL version of UNITREE. The ACL system attached to an IBM RS/6000 which is in turn connected to a HIPPI network.

Version 2:

The physical interface to the ACL is Ethernet. The logical interfaces uses the open system interconnection (OSI) model as is shown in Figure 2. RPC over TCP/IP is used on the session layer. The version of RPC used is SUN RPC version 2. On the Application Layer, SCSI-2 Cartridge Changer commands are used. This interface is described in Ampex document PD-19920918-1, NetSCSI specification. Over this interface a cartridge is accessed by a SCSI-2 Move Cartridge command specifying a source and destination address. The cartridge is identified by its location.

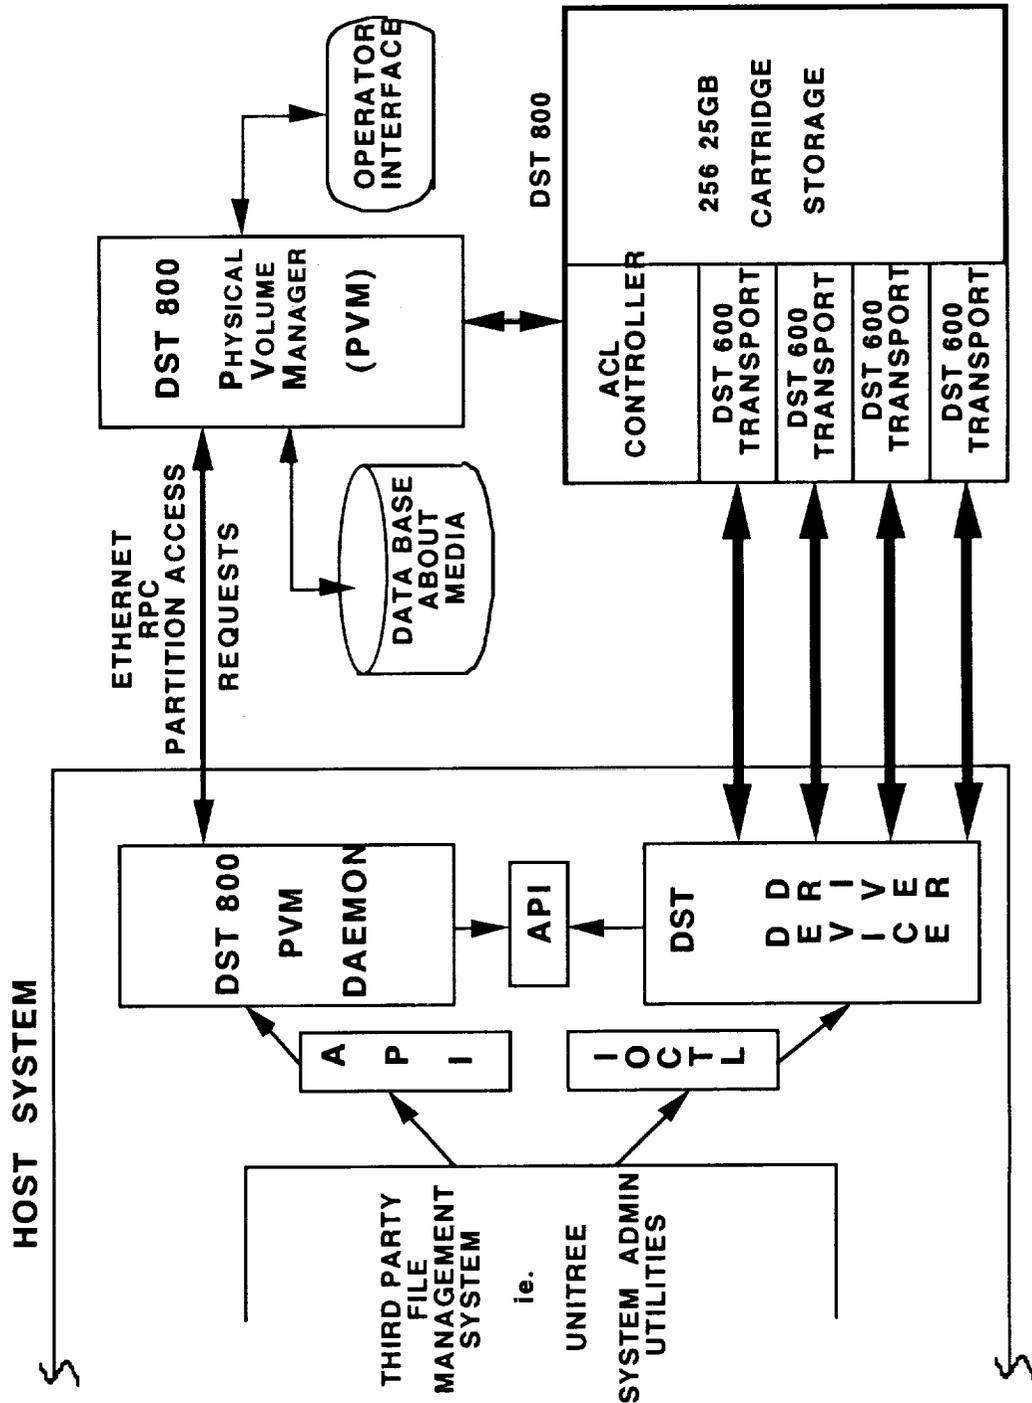


Figure 1. DST 800 With Tape Management Software

Whenever a cartridge that is not known to the DST 800 database is entered into the system for information exchange, the following procedure is performed.

1. The information contained in the cartridge manifest describing the attributes of the cartridge is entered in the DST 800 database.
2. A DST 800 barcode label is placed on the cartridge.
3. An operator command is issued to enter the cartridge into the DST 800.
4. The operator places the cartridge in an Import/Export slot.

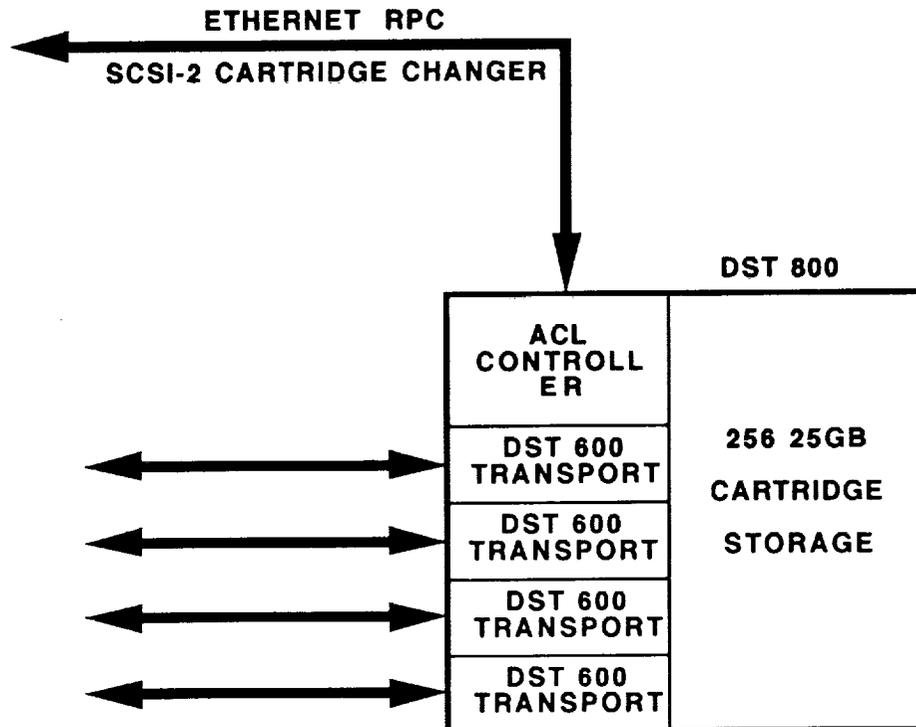


Figure 2. DST 800 Without Tape Management Software

The ACL automatically notifies the Physical Volume Manager (PVM) whenever a cartridge is introduced to one of its 8 Import/Export slots. The PVM automatically commands the ACL to read the barcode after receiving this notification. However, there is no requirement for the PVM to prioritize reading the barcode of a newly introduced cartridge over other operations. Barcodes are not automatically read when cartridges are loaded into tape drives. The ACL and PVM both maintain state information which includes a map of all cartridge locations (storage slots, import/export slots, tape drives) and, for all occupied locations, the barcode values (6 numeric characters). The barcode is an interleaved two out of five code format (ANSI standard MH 10.8-1983) and is a high density, self checking code with parity bits and accommodates bi-directional reading. The data is a six digit number with low optical density for reliable reading. The narrow bars are 30 mils and the wide bars are 75 mils.

Each cartridge is identified by a human readable cartridge ID on one end of the cartridge and a barcode on the other end. The IDs are placed on the short sides of the cartridge. The barcode is read horizontally from left to right and right to left. The DST library drive is designed to allow loading from the front of the ACL by the operator and at the back of the drive by the cartridge handling system. In the event of cartridge handling system failure, the PVM will communicate with the operator via a graphical interface on the front of the library and the operator can retrieve cartridges from the storage bins and mount them in the library drives. When automatic operation is restored, the ACL system will read the barcode on each cartridge to re-establish the integrity of the cartridge positions. This action takes just over 32 seconds.

All power supplies are equipped with over temperature sensors, that turn a particular supply off if an over temperature exists within that supply. In addition cooling fans are equipped with motion sensors that are monitored by the system. If a fan quits working, the system will shut down the subsystem affected. The environmental specifications for the ACL are the same as specified under Media Usage Life for operational environments.

Eight Import/Export slots can be used to input cartridges into the system or export cartridges out of the system. The DST 800 doors are only opened for repair or to remove cartridges for

manually loading of tape drives when the cartridge accessor is down for repair. The cartridge accessor always returns the cartridge to its original position.

The maximum number of drives that have been coupled with the ACL is four. In the television marketplace over 100 ACLs operate in the maximum configuration. The original design goal for the cartridge accessor was to support 4 transports, continuously playing back to back 7 second commercial spot segments. Today's typical cartridge accessor cycle times are summarized below for back to back operation.

	Min.	Ave	Max.	Description
Move	0.5	0.5	.05	Go to a drive
Picker (in/out)	1	1	1	Get a cartridge
Move	0.5	0.7	.09	Go to a bin
Picker (in/out)	1	1	1	Put cartridge
Move	0.5	0.7	.09	To another bin
Picker (in/out)	1	1	1	Get cartridge
Move	0.5	0.7	.09	Go to a drive
Picker (in/out)	1	1	1	Put cartridge
	6 sec	6.6 sec	7.2 sec	

© 1993 by Ampex Systems Corporation

